Alejandro Bernales – Nicolás Garrido – Satchit Sagade – Marcela Valenzuela – Christian Westheide

A Tale of One Exchange and Two Order Books: Effects of Fragmentation in the Absence of Competition

A Tale of One Exchange and Two Order Books: Effects of Fragmentation in the Absence of Competition

Alejandro Bernales   Nicolás Garrido   Satchit Sagade
Marcela Valenzuela   Christian Westheide*

First version: August 2016 - This version: October 2018

Abstract

Exchanges nowadays routinely operate multiple, almost identically structured limit order markets for the same security. We study the effects of such fragmentation on market performance using a dynamic model where agents trade strategically across two identically-organized limit order books. We show that fragmented markets, in equilibrium, offer higher welfare to intermediaries at the expense of investors with intrinsic trading motives, and lower liquidity than consolidated markets. Consistent with our theory, we document improvements in liquidity and lower profits for liquidity providers when Euronext, in 2009, consolidated its order flow for stocks traded across two country-specific and identically-organized order books into a single order book. Our results suggest that competition in market design, not fragmentation, drives previously documented improvements in market quality when new trading venues emerge; in the absence of such competition, market fragmentation is harmful.

Keywords: Fragmentation, Competition, Liquidity, Price Efficiency

JEL Classification: G10, G12

When you split these liquidity pools [...] what happens is that overall volumes tend to go up because the market starts to arbitrage and tries to put the market back together, the value of data goes up. And the whole thing for us turns out to be very good business [...] we don't think it's in the best interest of the market [...]

– Jeffrey Sprecher, Chairman and CEO, Intercontinental Exchange during the

Q1 2017 Earnings Call dated 03 May 2017

1. Introduction

In recent years, equity markets in the United States, the European Union, and elsewhere have evolved from national/regional stock exchanges being the dominant liquidity pools to a fragmented multi-market environment where the same stock now trades on multiple competing exchanges. These markets have simultaneously also experienced a process of consolidation as a result of national and international mergers between exchanges such that only a small number of operators, each running several exchanges, now compete with one another. For example, in the United States, the three largest exchange operators – Intercontinental Exchange, Nasdaq OMX, and Cboe – currently operate a total of eleven lit equity exchanges.[1] In most cases, the individual trading venues operated by a single operator employ almost identical rules and use the same technology such that differences between venues are minimal. This raises the question as to the effects of fragmentation when competition between venues is absent.

We examine the effects of fragmentation on market performance through a dynamic equilibrium model which characterizes such a multi-market environment. Our model builds on the single market models of Goettler et al. (2005, 2009). It is set up as a stochastic

---

[1]For ease of exposition, we use the terms *order book* and *venue* interchangeably while referring to an individual limit order market, and the terms *operator* and *exchange* interchangeably while referring to the venue's owner.

trading game in which a single asset can be traded in two identically-organized limit order books by agents who are heterogeneous in terms of their intrinsic economic reasons to trade the asset. They enter the market following a Poisson process and make optimal endogenous trading decisions depending on their private values, the state of both limit order books, the stochastically evolving fundamental value of the asset, expected costs of delaying order execution, and the possibility of reentering the market to revise existing limit orders. Limit orders in both order books are independently executed based on price and time priority. We parameterize the model following Goettler et al. (2009) and solve it numerically. By comparing a multi-market environment to a consolidated market setup, we analyze the effects of fragmentation across multiple venues not actively competing with each other.

Agents endogenously decide whether to provide or consume liquidity on a limit order book of their choice and in the presence of discrete prices. Those with an intrinsic motive to trade balance the delay costs associated with submitting limit orders and the immediacy costs associated with submitting market orders. Consequently, agents with large absolute private values are more likely to submit market orders. Agents with no intrinsic trading motives generate their profits solely from the trading process and are therefore more patient. Hence, they prefer to act as intermediaries by either submitting new limit orders or picking off existing limit orders that are mispriced.

We find that in a fragmented market, agents who provide liquidity submit less aggressive limit orders than in a consolidated market because they can circumvent time priority in one order book by submitting an order to the second order book. As this form of queue-jumping is impossible in a single venue, agents submit more aggressively priced limit orders. This reduction in price competition among liquidity providers in a fragmented market translates into lower price efficiency as well as higher quoted, effective, and realized spreads. Consequently, agents with (without) an intrinsic trading motive extract lower (higher) welfare gains in a fragmented market. Nevertheless, aggregate welfare does not differ markedly between a consolidated and fragmented market.

2

In an alternative parameterization, we solve the model in a fragmented market setting where the number of agents with zero private values is doubled. This allows us to approximate the effects of higher profits in fragmented markets leading to increased participation by intermediaries.[2] We find that higher competition between intermediaries leads to narrower bid-ask spreads as compared to a fragmented market under the baseline parameterization. However, the higher picking-off risk they generate leads agents who trade for intrinsic reasons to choose market orders more frequently. As these agents' limit orders are aggressively priced in the first place, the bid-ask spread in this setting is still inferior to that in a consolidated market under the baseline parameterization. The overall effect is an incremental shift in welfare towards intermediaries but no significant change in aggregate welfare.

We empirically test the model predictions by examining a unique event in which Euronext implemented a single order book per asset for their Paris, Amsterdam, and Brussels markets on 14 January 2009. It previously operated multiple independent order books for stocks cross-listed on these markets. This event led to a decrease in fragmentation for the affected stocks. We employ difference-in-differences estimations using a matched sample of treatment and control firms. Previous empirical studies examining the effects of new exchange operators entering a market, such as Foucault and Menkveld (2008), Hengelbrock and Theissen (2009) and Chlistalla and Lutat (2011), can be viewed as testing the joint effects of fragmentation and competition. This is because the entry of a new exchange, in addition to increasing fragmentation, also materially alters the competitive environment between operators. The new operator typically attempts to differentiate its platform along critical features such as trading speed, transaction fees, or the ability to execute large blocks. In contrast, the multiple order books operated by Euronext before the implementation of a single order book had identical trading protocols.

The empirical analysis broadly confirms the theoretical results. We find quoted spreads

---

[2]Ideally, one would incorporate endogenous participation in the presence of fixed entry costs in the model, which is, computationally challenging.

in the consolidated market to be lower by 35% than local spreads in the fragmented market before the event. Consistent with our model, effective spreads, both measured using local and inside quotes, are smaller after the implementation of the single order book. Higher competition post consolidation reduces the potential rent extraction by liquidity providers, resulting in at least 49% lower realized spreads. We also observe a weakly significant improvement in price efficiency.

While we are unable to empirically compute welfare effects, we find that the introduction of a single order book does not decrease the trading volume. This is despite the elimination of arbitrage trades between the multiple Euronext markets, which are responsible for up to 6.2% of the trading volume before the introduction of a single order book. This suggests that reduced transaction costs allow more participation by investors with intrinsic motives to trade implying an increase in their welfare, which is consistent with our theoretical results.

Our results contribute to the literature on equity market fragmentation.[3] Our contribution to this literature is twofold: First, we provide a dynamic multi-period model of multiple limit order markets that incorporates several real-world features and allows for more flexible agent behavior as compared to previous models. We allow competition between heterogeneous agents who arrive sequentially based on an exogenous arrival rate and make endogenous trading decisions such as where to submit an order, whether to submit limit or market orders, what limit price to choose in case of a limit order, and whether/how to modify standing limit orders. The existence of real world frictions such as price discreteness, the absence of perfect competition as argued for by Glosten (1998), and endogenous liquidity provision drives our results. Second, by analyzing a unique event that isolates fragmentation from competition between venues, we provide empirical evidence that fragmented markets are liquidity- and likely welfare- reducing, which is consistent with our theoretical results.

Early theories on fragmentation such as Mendelson (1987), Pagano (1989), Chowdhry and Nanda (1991) highlight the positive network externalities generated by consolidating

---

[3]See Gomber et al. (2017) for a detailed survey of this literature.

trading on a single venue. However, a consolidated market is no longer the equilibrium outcome in the absence of post-trade transparency (Madhavan, 1995) and in the presence of real world frictions such as differences in markets' absorptive capacity and institutional mechanisms (Pagano, 1989), order splitting behavior (Chowdhry and Nanda, 1991), or trader heterogeneity (Harris, 1993).

Parlour and Seppi (2003) analyze market fragmentation in a richer model containing institutional details such as a discrete pricing grid and a specialist market competing with a pure limit order book. They show that these markets can coexist in equilibrium and that the effects on liquidity are ambiguous. In a study closely related to modern equity markets, Foucault and Menkveld (2008) model fee-based competition between two operators and predict that the entry of a second exchange increases consolidated depth, and that increased use of smart order routers leads to an increase in liquidity in the entrant market. In contrast, our model allows for endogenous choice between liquidity supply and consumption and does not require positive order submission fees.

Fragmented equity markets typically induce a downward pressure on trading fees but also lead to increased adverse selection due to cross-market arbitrage. Baldauf and Mollner (2018) show that the net effect of these two opposing forces is context-dependent. Trading venues also discriminate between the fees charged to suppliers and consumers of liquidity. For instance, many markets provide rebates to liquidity providers (maker-taker pricing) or consumers (taker-maker pricing). Recent theories investigating the effect of maker-taker pricing on market quality show that if prices are continuous, only the net fees are relevant (Colliard and Foucault, 2012), whereas in the presence of discrete prices the fee breakdown affects agents' order submission behavior (Foucault et al., 2013). Panayides et al. (2017) extend this insight to fragmented markets and show that maker-taker fees lead to a migration of limit and market orders across venues. Chao et al. (2018) show that, in the presence of a fixed tick size, exchanges use such alternative fee structures to attract traders to their market by offering a finer (implicit) pricing grid. In our model, we do not rely on trading

5

fees for a fragmented market to coexist in equilibrium.

Several studies document improvements in market outcomes when rival operators begin competing with incumbent trading venues (Boehmer and Boehmer, 2003; Chlistalla and Lutat, 2011; Hengelbrock and Theissen, 2009; Nguyen et al., 2007). Haslag and Ringgenberg (2015) document differences in the impact of fragmentation across stocks whereas Chung and Chuwonganant (2012) find that fragmentation leads to a general deterioration in market quality. Degryse et al. (2015) and Gresse (Gresse) differentiate between lit and dark fragmentation and find that the former improves liquidity. They, however, disagree on the effects of dark fragmentation. These studies investigate the joint effects of fragmentation and competition among trading venues. Amihud et al. (2003) study a reduction in fragmentation on the Tel Aviv Stock Exchange resulting from the exercise of deep in-the-money share warrants and, consistent with our results, find an increase in liquidity. However, the stocks and warrants in their sample were not perfectly fungible assets and, differently from modern equity markets, were traded in single or periodic batch auctions. In contrast to this literature, our setting features within-operator fragmentation of trading for the same asset.

The shift in welfare towards intermediaries, especially when the amount of intermediation increases, has strong regulatory implications. Higher revenues earned by intermediaries should lead to excessive investment in their capacities and, in the absence of a meaningful increase in aggregate welfare, a socially wasteful arms race. From a regulatory standpoint, this raises the question whether restricting within-operator fragmentation would lead to improvements in social welfare.

The remainder of the paper is structured as follows. Section 2 describes the theoretical model. In Section 3, we provide a partial equilibrium analysis to explain the intuition behind agents' order submission strategies. Section 4 compares the model implications for trader behavior, market quality, and welfare in consolidated and fragmented markets. In Section 5, we present the empirical analysis of Euronext's introduction of a single order book. Finally, we conclude in Section 6.

6

2. Multi-Market Model

2.1 Model Setting

We consider an economy in continuous time with a single financial asset that trades on two independent and identically organized limit order books. The trading activity reflects a sequential non-cooperative game, where agents make endogenous, asynchronous decisions to maximize their expected payoffs, taking into account their private reasons to trade the asset, market conditions, and the strategies employed by agents expected to arrive in the future.

Each limit order book at time t and in market m with $m \in \{1, 2\}$, $L_{m,t}$, is characterized by a set of discrete prices denoted by $\{p_m^i\}_{i=-N}^N$, where $p_m^i < p_m^{i+1}$ and N is a finite number. Let d be the distance between any two consecutive prices, which we refer to as the tick size (i.e., $d = p_m^{i+1} - p_m^i$). The tick size in both limit order books is equal. Let $l_{m,t}^i$ be the queue of unexecuted limit orders in order book m at time t and price p_m^i. A positive (negative) $l_{m,t}^i$ denotes the number of buy (sell) limit orders, and it represents the depth of the book $L_{m,t}$ at price p_m^i. In the book $L_{m,t}$ at time t, the best bid price is $B(L_{m,t}) = sup\{p_m^i | l_{m,t}^i > 0\}$ and the best ask price is $A(L_{m,t}) = inf\{p_m^i | l_{m,t}^i < 0\}$. $B(L_{m,t}) = -\infty$ or $A(L_{m,t}) = \infty$ if the order book $L_{m,t}$ is empty at time t on the buy side or on the sell side, respectively. Each limit order book independently respects price and time priority while executing the limit orders, i.e., buy (sell) limit orders at higher (lower) prices have priority in the queue and limit orders submitted earlier at the same price are executed first.

The fundamental value of the asset, v_t, is stochastic and known by the agents; its innovations follow an independent Poisson process with parameter λ_v. In case of an innovation, the fundamental value increases or decreases by one tick with equal probability. The economy is populated by risk-neutral agents who arrive sequentially following a Poisson process with intensity λ and trade one share before exiting. All agents observe both limit order books (i.e., prices and depths at each price) and can submit a limit or market order to either book. Agents can reenter the market to modify unexecuted limit orders following a Poisson pro-

7

cesses with parameter λ_r, which is the same for both markets and is independent of λ. The resulting lack of instantaneous reentry implies that decisions regarding limit order submissions are sticky. Agents face opportunity and monitoring cost when they cannot immediately trade the asset. This delay cost, denoted by a discount rate $\rho \in [0,1]$, is constant across agents and order books, and applies to their total payoff.

Agents are heterogeneous in terms of their intrinsic economic motives to trade the asset. These motives are reflected in their private values. Each agent has a private value, α, which is known to her. α is drawn from the discrete vector $\Psi = \{\alpha_1, \alpha_2, ..., \alpha_g\}$ using a discrete distribution, F_α, where g is a finite integer. Private values reflect the fact that agents would like to trade for various reasons unrelated to the fundamental value of the asset (e.g. hedging needs, tax exposures and/or wealth shocks). They are idiosyncratic and constant for each agent.

Agent heterogeneity, delay costs and the fundamental value of the asset determine agents' trading behavior. On the one hand, suppose agent x with a positive private value (i.e., $\alpha > 0$) arrives at time t_x. This agent is likely to be a buyer because she would like to have the asset to obtain the intrinsic benefit reflected by α. In this case, the agent's expected payoff is: $(\alpha + v_{t'} - p)e^{-\rho(t'-t_x)}$, where p is the transaction price, t' is the expected time of the transaction, and $v_{t'}$ is the expected fundamental value of the asset at time t'. Moreover, if α is very high, her delay cost, denoted by $(e^{-\rho(t'-t_x)} - 1)\alpha$, is correspondingly high and, hence, she may prefer to buy the asset as soon as possible by using a market order. In this case, the agent will pay an immediacy cost denoted by $(v_{t'} - p)^{-\rho(t'-t_x)}$ The agent will accept this immediacy cost because she is mainly generating her profits from the large private value, α, rather than from the trading process *per se*. In this sense, an agent with a high absolute private value will probably be a liquidity taker.[4]

On the other hand, suppose an agent y with a private value equal to zero (i.e., $\alpha = 0$)

---

[4]A similar example can be explained in the other direction in case of an agent with a negative private value (i.e., $\alpha < 0$) having a preference to sell.

arrives at time t_y. This agent needs to find a profitable opportunity purely in the trading process, by obtaining a good price relative to the fundamental value, because she does not obtain any intrinsic economic benefit from trading. Consequently, she may be patient and prefer to act as a liquidity provider, in turn receiving the immediacy cost paid by a liquidity taker. Alternatively, she may trade aggressively against a standing limit order that is mispriced relative to the fundamental value. Note that agents with $\alpha = 0$ are indifferent with respect to taking either side of the market because they can maximize their benefits by either selling or buying the asset.

Liquidity providers are exposed to the risk of being picked off because limit orders can generate a negative payoff if they are in an unfavorable position relative to the fundamental value. For example, suppose an agent x with $\alpha = 0$ first arrives at time $t = 0$ and submits a limit buy order to set the best bid price, B in market $m = 1$. Suppose further that at time t^*, the fundamental value of the asset decreases to level v_{t^*}, such that $v_{t^*} < B$ and simultaneously another agent, denoted y, with private value $\alpha = 0$ arrives. Since agent x cannot immediately modify her unexecuted limit order, agent y can submit a market sell order and pick off agent x's order generating an instantaneous profit equal to $(B - v_{t^*})$. Agent x, on the other hand, has a negative realized payoff given by $(v_{t^*} - B)e^{-\rho t^*}$.

We center each limit order book at the contemporaneous fundamental value of the asset, i.e., by setting $p_m^0 = v_t$. Suppose at time $t = 0$ the fundamental value is v_0, but after a period τ the fundamental value experiences some innovations and its new value is v_τ, with $v_\tau - v_0 = qd$, where q is a positive or negative integer. In this case, we shift both books by q ticks to center them at the new level of the fundamental value v_τ. Thus, we move the queues of existing limit orders in both books to take the relative difference with respect to the new fundamental value into account. This implies that prices of all orders are always relative to the current fundamental value of the asset. This transformation allows us to greatly reduce the dimensionality of the state-space because agents always make decisions in terms of prices

9

relative to the fundamental value.[5]

Each agent takes three main trading decisions upon arrival: i) to submit an order to $L_{1,t}$ or $L_{2,t}$; ii) to submit a buy or a sell order; and iii) to choose the limit price, which implies the decision to submit either a market or a limit order.[6] As mentioned above, an agent can re-enter the market and modify her unexecuted limit order. Hence, she has to take the following additional trading decisions after re-entering: i) to keep her unexecuted limit order unchanged or to cancel it; ii) in case of a cancellation, to submit a new order to $L_{1,t}$ or $L_{2,t}$; iii) to choose whether the new order will be a buy or a sell order; and iv) to choose the price of the new order. The decision to leave the order unchanged has the advantage of maintaining its time priority in the respective queue. The disadvantage is the exposure to picking-off or non-execution risk depending on the direction of change in fundamental value since the initial order submission.

2.2 Agents' Dynamic Maximization Problem and Equilibrium

There is a set of states $s \in \{1, 2, \ldots, S\}$ that describes the market conditions in the economy. These market conditions are observed by each agent before making any decision. The state s that an agent observes is described by the contemporaneous limit order books, L_1 and L_2; the agent's private value α; and in the case that the agent previously submitted a limit order to any of the books, the status of that order in L_1 or L_2, i.e., its original submission price, its queue priority in the book, and its type (i.e., buy or sell). The fundamental value of the asset, v, is implicitly part of the variables that describe the state s, since agents interpret

---

[5]Note that under this normalization, we can still observe limit orders being picked-off. For example, suppose that the current time is t and the fundamental value is v_t; hence $p_m^0 = v_t$. Suppose, that the current bid price is $B(L_{m,t}) = p_m^{-1}$ and the ask price is $A(L_{m,t}) = p_m^2$. Subsequently, at time t_{po}, if the fundamental value decreases by twice the amount of the tick size (i.e., $q = -2$), after centering the book, the bid and ask prices are $B(L_{m,t_{po}}) = p_m^1$ and $A(L_{m,t_{po}}) = p_m^4$, respectively. Thus, a newly arriving agent can submit a market sell order against the limit order at the bid price to generate a profit. Consequently, the limit order at p_m^1 will disappear, and the new bid price will be equal to the price at the center of the book (i.e., $B(L_{m,t_{po}+\Delta t}) = p_m^0$, where Δt is the time until the limit buy order above the fundamental value is picked-off).

[6]A buy (sell) limit order priced at or above (below) the best ask (bid) is equivalent to a market buy (sell) order.

limit order prices relative to the fundamental value. For convenience, we set the arrival time of an agent to zero in the following discussion.

Let $a \in \Theta(s)$ be the agent's potential trading decision, where $\Theta(s)$ is the set of all possible decisions that an agent can take in state s. Suppose that the optimal decision given state s is $\tilde{a} \in \Theta(s)$. Let $\eta(h|\tilde{a}, s)$ be the probability that the optimally submitted order is executed at time h. The probability $\eta(\cdot)$ depends on future states and potential optimal decisions taken by other agents up to time h. If the agent submits a market order, then the probability $\eta(0|\tilde{a}, s)$ is equal to one, while $\eta(h|\tilde{a}, s)$ converges to zero if the agent submits a limit order further away from the fundamental value of the asset. Let $\gamma(v|h)$ be the density function of v at time h, which is exogenous and characterized by the Poisson process of the fundamental value of the asset with rate λ_v. Thus, the expected value of the optimal order submission $\tilde{a} \in \Theta(s)$, if the order is executed prior to the agent's re-entry time h_r, is:

$$\pi(h_r, \tilde{a}, s) = \int_0^{h_r} \int_{-\infty}^{\infty} e^{-\rho h} \left((\alpha + v_h - \tilde{p})\tilde{x} \right) \cdot \gamma(v_h|h) \cdot \eta(h|\tilde{a}, s) dv_h dh \tag{1}$$

where $\tilde{p}$ is the submission price and $\tilde{x}$ is the order direction indicator (i.e., $\tilde{x} = 1$ if the agent buys and $\tilde{x} = -1$ if the agent sells) and both are components of the optimal decision $\tilde{a}$. The expression $(\alpha + v_h - \tilde{p})\tilde{x}$ is the instantaneous payoff, which is discounted back to the trader's arrival time at rate ρ.

Let $\psi(s_{h_r}|h_r, \tilde{a}, s)$ be the probability that state s_{h_r} is observed by the agent at her re-entry time h_r, given her decision $\tilde{a}$ taken in the previous state s. The probability $\psi(\cdot)$ depends on the states and potential optimal decisions taken by other agents up to time h_r. In addition, let $R(h_r)$ be the cumulative probability distribution of the agent's re-entry time, which is exogenous and described by the Poisson process λ_r. Thus, the Bellman equation that describes the agent's problem of maximizing her total expected value, $V(s)$, after arriving in state s is given by:

$$V(s) = \max_{\tilde{a} \in \Theta(s)} \int_0^{\infty} \left[\pi(h_r, \tilde{a}, s) + e^{-\rho h_r} \int_{s_{h_r} \in S} V(s_{h_r}) \cdot \psi(s_{h_r}|h_r, \tilde{a}, s) ds_{h_r} \right] dR(h_r) \tag{2}$$

where S is the set of possible states. The first term is defined in Equation (1), and the second term describes the subsequent payoffs in the case of re-entries.

In equilibrium, each agent behaves optimally by maximizing her expected utility, based on the observed state that describes market conditions (as in Equation (2)). In this sense, optimal decisions are state dependent. They are also Markovian, because the state observed by an agent is a consequence of the previous states and the historical optimal decisions taken in the trading game. As there is competition between agents, the equilibrium is competitive (although there is no competition between venues). A competitive equilibrium means that competing agents would respond to an agent's local deviation in a way that leads to a reduction in the deviating agent's expected utility. We obtain a stationary and symmetric equilibrium, as in Doraszelski and Pakes (2007). In such an equilibrium, optimal decisions are time independent, i.e., they stay the same when an agent faces the same state in the present or in the future.

The trading game is also Bayesian in the sense that an agent knows her intrinsic private value to trade α, but she does not know the private values of other agents that are part of the game. Hence, our solution concept is a Markov perfect Bayesian Equilibrium (see Maskin and Tirole, 2001). In the trading game, there is a state transition process where the probability of arriving in state s_{h_r} from state s is given by $\psi(s_{h_r}|\tilde{a}, s, h_r)$. Thus, two conditions must hold in the equilibrium: agents solve equation (2) in each state s, and the equilibrium beliefs are consistent for each state over time.

As mentioned earlier, the state s is defined by the four-tuple ($L_{1,t}$, $L_{2,t}$, α, *status of previous limit order*), where all variables that describe the state are discrete. Moreover, each agent's potential decision a is taken from $\Theta(s)$, which is the set of all possible decisions that can be taken in state s. This set of possible decisions is discrete and finite given the features of the model. Consequently, the state space is countable and the decision space is finite; thus the trading game has a Markov perfect equilibrium (see Rieder, 1979). Despite the fact that the model does not lend itself to a closed-form solution, we check whether the equilibrium

12

is computationally unique by using different initial values.

2.3 Solution Approach and Model Parameterization

Given the large dimension of the state space, we use the Pakes and McGuire (2001) algorithm to compute a stationary and symmetric Markov-perfect equilibrium. The intuition behind this algorithm is that the trading game by itself can be used initially as a tool in which agents learn how to behave in each state. Thus, we set the initial beliefs about the expected payoffs of potential decisions for each state. Agents take the trading decision that provides the highest expected payoff conditional on the state they observe. Subsequently, agents dynamically update their beliefs by playing the game and observing the realized payoffs of their trading decisions. In this sense, the algorithm is based on agents following a learning-by-doing mechanism. The Pakes and McGuire (2001) algorithm is able to deal with a large state space because it reaches the equilibrium only on the recurring states class.

The equilibrium is reached when there is nothing left to learn, i.e., when beliefs about expected payoffs have converged. We apply the same procedure used by Goettler et al. (2009) to determine whether the equilibrium is reached. Once we reach the equilibrium after making the agents play the game for at least 10 billion trading events, we fix the agents' beliefs and simulate a further 20 million events. All theoretical results presented in Section 4 are computed from the latter.

We parameterize the model following Goettler et al. (2009). We set the intensity of the Poisson process followed by agents' arrival λ to one. The intensity of the Poisson process followed by agents' re-entry λ_r is set to 0.25; the intensity of the Poisson process followed by the innovations of the fundamental value λ_v is set to 0.125. We set the tick size d in both order books to one, and the number of discrete prices available on each side of both order books to $N = 31$. The delay cost ρ is set to 0.05. The private value α is drawn from the discrete vector $\Psi = \{-8, -4, 0, 4, 8\}$ using the cumulative probability distribution

13

$F_\alpha = \{0.15, 0.35, 0.65, 0.85, 1.0\}$.

While market entry is exogenous in our model, we posit that, if entry were costly and endogenous, the composition of agents' population would change. In particular, higher trading profits generated by agents of type $\alpha = 0$ in fragmented markets would increase their participation.[7] In a computationally simpler alternative, we create an additional parameter configuration by keeping the arrival rates of agents with non-zero private value unchanged and doubling the arrival rate of agents with private value equal to zero. In other words, we set the intensity of agent arrival to 1.3 and draw the different agent types from the cumulative distribution $F_\alpha = \{0.15/1.3, 0.35/1.3, 0.95/1.3, 1.15/1.3, 1.0\}$.

As a robustness check, we run the model with different parameter setups. We multiply the following original parameters from Goettler et al. (2009) by 0.8 and 1.2: the cost of delay, ρ; the agents' arrival intensity λ; the innovation arrival intensity of the fundamental value, λ_v; and agents' re-entering intensity λ_r. We also modify the distribution of agents type in terms of their private values by running a setup with a subgroup of agent types $\Psi = \{-4, 0, 4\}$ using the cumulative probability distribution $F_\alpha = \{0.35, 0.65, 1.0\}$. Further, we run a setup where the volatility of the fundamental value is set to zero. Finally, we also run a setup with a higher tick size. The results obtained from all these three robustness checks are qualitatively similar to the results presented here.

3. Partial Equilibrium Analysis: Price Competition

In this section, we explain the different agents' trading strategies in a single and fragmented market in a partial equilibrium. For simplicity, and only in this section, we assume that the fundamental value does not change over time (i.e., $v_t = v$) and that there is a single agent type with private value α^* who wants to buy one share of the asset. Agents arrive following a Poisson process with intensity λ^*.

---

[7]We would expect a decrease in participation for agents with non-zero private values if their expected payoffs were sufficiently small relative to the entry cost. In what follows, we assume that it is always optimal for agents with non-zero private values to participate.

First, consider a single market setup, with an exogenously determined bid price B submitted by Trader 0, and an empty sell side of the book (i.e., the ask price is infinite). Limit orders are executed based on price and time priority. To keep this partial equilibrium simple, we assume that agents can instantaneously modify unexecuted limit orders, they know that market sell orders arrive following a Poisson process with intensity λ^{MO}, and the tick size is very small.

A new trader, Trader 1, enters the market and prefers to submit a limit buy order. Hence, Trader 0 and Trader 1 will compete to execute their limit orders against incoming market sell orders. If agents were not able to modify their limit orders, Trader 1 would submit a limit order at a slightly higher price $B' = B + \Delta$, where Δ is very small. In this case, the expected payoff of Trader 1 would be strictly larger than the expected payoff of Trader 0. However, Trader 0 can modify her limit order and would do so if Trader 1 followed this strategy, resulting in an escalation of price improvements until zero expected profit is obtained.

Nevertheless, Trader 1 can select price $p_{SM}^B > B$ to avoid these price escalations. We equate the expected payoffs of both traders to find this equilibrium price p_{SM}^B. Thus, in a competitive equilibrium, we have:

$$e^{-\rho \cdot h_{SM}^0} \left[\alpha^* + v - B \right] = e^{-\rho \cdot h_{SM}^1} \left[\alpha^* + v - p_{SM}^B \right], \tag{3}$$

where h_{SM}^0 (h_{SM}^1) is the expected execution time of the buy limit order at price B (p_{SM}^B) for Trader 0 (1). The left hand side of Equation (3) reflects Trader 0's expected payoff, while the right hand side of this equation represents Trader 1's expected payoff, where $h_{SM}^0 > h_{SM}^1$ and $p_{SM}^B < \alpha^* + v$. Then, p_{SM}^B is:

$$p_{SM}^B = B + (1 - e^{-\rho \cdot (h_{SM}^0 - h_{SM}^1)}) \left[\alpha^* + v - B \right] \tag{4}$$

Thus, in this partial equilibrium, Trader 0 submits a buy limit order at price B, while Trader 1 submits a buy limit order at price p_{SM}^B. Both traders have the same expected payoff. Trader 0 buys at a lower price than Trader 1, but Trader 0 has to wait longer for the

15

execution of her limit order than Trader 1. These effects cancel out such that the expected payoffs are equal for both traders and p_{SM}^B is the new best bid price of the market.

The competitive equilibrium is very different in a fragmented market setting. Let us now consider the same market features used to obtain Equation (4), but this time Trader 1 enters a fragmented market with two identical limit order books. When she enters there is one buy limit order at price B in one limit order book, and there are no limit sell orders. The limit order at price B was submitted by Trader 0. In this setting, traders can submit orders to any of the two order books.

The trader needs to decide between submitting a limit buy order at the bid price B or a more aggressive limit buy order at price $p_{MM}^B > B$ in any of the two books. As in the single market setting, Trader 1 selects her bid price such that in equilibrium both traders have the same expected payoff, which can be written as:

$$e^{-\rho \cdot h_{MM}^0} \left[\alpha^* + v - B \right] = e^{-\rho \cdot h_{MM}^1} \left[\alpha^* + v - p_{MM}^B \right], \tag{5}$$

where h_{MM}^0 (h_{MM}^1) is the expected execution time of the buy limit order at price B (p_{MM}^B) for Trader 0 (1) in a multi-market setup, with $p_{MM}^B < \alpha^* + v$. However, since there are two limit order books, Trader 1 can also submit a buy limit order at price B in the competing market, which has the same expected execution time as the buy limit order submitted by Trader 0 (i.e., $h_{MM}^0 = h_{MM}^1$). Then, we have:

$$p_{MM}^B = B \tag{6}$$

Figure 1 shows Traders 1's expected payoff as a function of the buy limit order price under a single and fragmented market setup. This figure shows that the slope of Trader 1's expected payoffs is $e^{-\rho \cdot h_{SM}^1}$ and $e^{-\rho \cdot h_{MM}^1}$ under the single and fragmented market setup, respectively. We can observe that $e^{-\rho \cdot h_{SM}^1} > e^{-\rho \cdot h_{MM}^1}$ due to Trader 1's queue-jumping behavior in the fragmented market setup described above. This is because a market order arriving in a fragmented setting can be executed against the limit orders of either Trader 0

16

or Trader 1, as these orders are submitted at the same price. Conversely, in a single market, Trader 1 always executes her limit order before Trader 0; thus $h^1_{SM} < h^1_{MM}$.

Figure 1 about here

The queue-jumping feature of the fragmented market also affects Trader 0's expected execution time, but in the opposite direction (see Figure 1). In a single market her order always sits behind Trader 1's limit order. However, in a fragmented market both traders' orders sit at the same price B, with an equal probability of execution. This reduces Trader 0's expected order execution time in a fragmented market as compared to a single market (i.e., $h^0_{SM} > h^0_{MM}$). Thus, Trader 0's expected payoff is larger in a fragmented market than in a single market.

It is important to note that there is a second potential competitive equilibrium in the fragmented market setup, but it is always suboptimal. In this equilibrium, Trader 1 can submit a buy limit order at a price $p^{B\dagger}_{MM} > B$ such as $e^{-\rho \cdot h^{0\dagger}_{MM}} [\alpha^* + v - B] = e^{-\rho \cdot h^{1\dagger}_{MM}} \left[\alpha^* + v - p^{B\dagger}_{MM}\right]$. In other words, Trader 1 mimics the trading strategy implemented in a single market, despite being in a fragmented market setup. In this equilibrium, Trader 0 has a lower price than Trader 1, but Trader 0 has to wait longer for the execution of her limit order than Trader 1 (as in the single market partial equilibrium).

Nevertheless, this second equilibrium would reduce the potential expected payoff for both traders. This is because if Trader 1 submits her order at $p^{B\dagger}_{MM} > B$, Trader 0's execution time increases (Trader 0 has to wait for Trader 1's order to execute first) relative to the case where Trader 1 submits an order at B in the second book (where both traders have the same probability of executing their orders). In other words, the parallel gray line of Figure 1 (which represents the expected payoff of Trader 0 in a multi-market) shifts downwards. This is because $h^{0\dagger}_{MM} > h^0_{MM}$ implies that $e^{-\rho \cdot h^{0\dagger}_{MM}} [\alpha^* + v - B] < e^{-\rho \cdot h^0_{MM}} [\alpha^* + v - B]$. Furthermore, we know that, in equilibrium, Trader 1's expected payoff 1 should be equal to Trader 0's expected payoff. Thus, Trader 1 would also obtain Trader 0's lower expected

17

payoff in this equilibrium as compared to the alternative strategy where she submits a limit order in the second market at price B. Thus, Trader 1 never submits an order at price $p_{MM}^{B\dagger} > B$.

The results of Equation (4) and (6) help explain the differences between the single and fragmented market setting. Equation (6) shows that in a fragmented market partial equilibrium with two traders, the bid price does not increase upon Trader 1's arrival. In other words, agents are less competitive in terms of their order aggressiveness in a fragmented market as compared to a single market. In the single market setup described in equation (4), Trader 1 can reduce the execution time of her buy order only by jumping ahead of Trader 0's limit order. In other words, queue-jumping is only possible by improving upon Trader 0's limit price (i.e., by submitting a more aggressive order) as limit orders obey price and time priority. Under this competitive environment, Trader 1 can submit a limit order at an equilibrium price $p_{SM}^{B} > B$, where both traders have the same expected payoff. Thus, price competition between traders increases the bid price in a single market setting. In a fragmented market setup, Trader 1 can jump ahead of the existing limit order at price B (submitted by Trader 0) not only by submitting a more aggressive limit order, but also by matching Trader 0's price B in the second market. This form of queue-jumping in a fragmented market reduces limit orders aggressiveness as compared to a single market.

4. Theoretical Implications

We examine the theoretical implications of market fragmentation for trader behavior, market quality, and welfare. To this end, we generate a simulated dataset under the following three market setups: i) a consolidated market with one limit order book; ii) a fragmented market with two identical limit order books; and iii) a fragmented markets with two identical limit order books and after doubling the population of agents with zero private values. We compute

mean levels of the variables of interest under these three market settings.[8]

4.1 Trading Behavior

Agents' order submission strategies determine the liquidity characteristics and the price discovery process of the asset, and, as a consequence, determine the welfare of different agent types and that of the economy. Hence, we start by analyzing agents' trading behavior in the three settings.

Table 1 presents the distribution of executed limit and market orders by each agent type and the probability of a new limit order being placed at the best quotes. Consistent with the intuition outlined in Section 2.1, agents with $\alpha = 0$ ($\alpha = 8$) execute 77.1% (80.4%) of their trades using limit (market) orders in the single market setting. Agents with $\alpha = 4$ use limit and market orders roughly equally. These frequencies remain largely unchanged in a fragmented market setting when the agents' population remains constant. Further, we observe that in a single market setting 50% of the limit orders are submitted at the best price, whereas this is the case for only 41.6% of limit orders in the fragmented market. This suggests that the overall degree of price competition is substantially higher in a consolidated market.

Table 1 about here

Under the alternative parametrization, agents with $\alpha = 0$ ($\alpha \neq 0$) further increase their reliance on limit (market) orders. Specifically, agents with $\alpha = 0$ execute 90% of their trades using limit orders and those with $\alpha = 4$ and $\alpha = 8$ execute 76.6% and 94.4% of their trades using market orders, respectively. Further, we find that the probability of submitting limit orders at the best quotes drops to 36%. This is because of an increase in the total frequency

---

[8]We do not report standard errors because a large number of trader arrivals leads to standard errors being sufficiently low such that a difference in means to the order of 10^{-2} is significantly different from zero.

19

of limit orders due to a combination of an increase in the population of $\alpha = 0$ agents and a higher probability of them using limit orders.[9]

We next compare limit order time to execution. We define time to execution as the difference between the time an order executes and the market entry time of the agent submitting the limit order. Panel A of Table 2 reports the results. Under the baseline parameterization, intermediaries' average time to execution decreases from 14.9 units in a consolidated market to 11.5 units in a fragmented market because the need to reprice existing limit orders is reduced when queues are shorter. For the other agent types there is little change. If we double the population of $\alpha = 0$ agents, the time to execution is higher (15.4 units) for intermediaries but lower for agents with an intrinsic trading motive. For the former this is clearly due to increased competition with other similar agents. For the latter, conditional on them quoting aggressively priced limit orders, counterparties arrive faster because of the higher population of $\alpha = 0$ agents.

Table 2 about here

Next, we examine the price setting behavior of executed limit orders for the different agent types by comparing the limit order price with the fundamental value at the time of submission. Panel B of Table 2 contains the results. Under all parameterizations, agents with $\alpha = 8$ submit limit orders with a negative difference between the limit price and the fundamental value of almost one tick in order to attract market orders and avoid excessive delay costs. Conversely, $\alpha = 0$ agents submit limit orders with a positive difference to earn the bid-ask spread and avoid the risk of immediately being picked off. Under the baseline parameterization, agents' ability to circumvent time priority by queue jumping is an important determinant of the level of price competition. In a single market, as this form of queue jumping is impossible, all agents submit limit orders that are, on average,

---

[9]In untabulated results, we observe that, under the alternative parametrization, the absolute frequency of orders submitted at or improving the best quotes is higher (lower) than a fragmented (single) market under the baseline parametrization.

slightly more aggressively priced as compared to a fragmented market (for example, 1.03 versus 1.06 ticks for $\alpha = 0$ agents). When we double the number of intermediaries, $\alpha = 8$ agent types submit limit orders that are priced most aggressively as compared to their price setting behavior under the two original parameterizations. The limit orders of $\alpha = 0$ agent types are priced more aggressively as compared to a fragmented market under the baseline parameterization but not as aggressively as the single market.

More aggressive limit orders in a single market compared to a fragmented market setting with the same arrival rates across agent types lead to higher ex ante picking-off risk. This should induce agents to more frequently cancel their orders to manage this higher risk. The share of executed limit orders actually picked off should reflect the net effect of these two forces. Panel C of Table 2 reports these results. We find that, in aggregate, the picking-off risk declines from 21.8% in a single market to 20.8% in a fragmented market under the baseline parametrization. However, these aggregates substantially mask the large differences across the agent types. Agents with $\alpha = 8$ ($\alpha = 0$) have the highest (lowest) probability of being picked off. In line with lower limit order aggressiveness in a fragmented market (see Panel B of Table 2), the picking-off risk for agents with $\alpha = 8$ ($\alpha = 0$) decreases from 73.9% (4.1%) to 72.1% (3.0%). The picking-off risk for all agent types is the highest when the arrival rate of intermediaries is doubled. This is due to a combination of higher price competition relative to the originally parameterized fragmented market and a higher arrival rate of agents who are more likely to pick off standing limit orders. However, overall picking-off risk decreases due to a change in the composition of agent types submitting limit orders.

4.2 Market Quality

In this subsection, we compare consolidated and fragmented markets in terms of the major market quality characteristics, i.e., liquidity and price efficiency. We compute all market quality measures based on *local quotes* and *inside quotes*. The former comprise the bid and

21

ask prices in one of the two markets, whereas the latter comprise the highest bid and the lowest ask across the two limit order books. The two are obviously identical in a single market setting. Table 3 provides the results.

Table 3 about here

We first measure daily quoted liquidity by computing time-weighted quoted bid-ask spreads. We also compute the time-weighted market depth, i.e., the total number of limit orders waiting to be executed at the top of the book. Panel A of Table 3 reports the results. We find that fragmentation impairs quoted liquidity when the population of agents remains unchanged. This is indicated by wider spreads and lower depth in a market with two limit order books. In particular, local and inside quoted spreads are lower by approximately 1.0 and 0.3 ticks, respectively, in a single market setting. In fact, bid-ask spreads in a consolidated market are the lowest across all three parameterizations, although the difference compared to a fragmented market with a double arrival rate of $\alpha = 0$ agents is slightly smaller. This result is consistent with higher price competition between liquidity providers in a single market, as explained in Sections 3 and 4.1.

Local time-weighted top-of-book depth is lower by more than 30% and inside depth is lower by almost 9% in a fragmented market as compared to the single market under the baseline parameterization. In a fragmented market, the total order flow is distributed across the two order books leading to lower local depth in each order book. The higher inside depth in the single market can be explained by the higher time to execution in this setting (see Section 4.2). Specifically, while the order arrival rate remains unchanged, higher price competition induces agents to more frequently revise existing orders leading to longer execution time and, in turn, higher time-weighted inside depth. The results change when we double the participation of $\alpha = 0$ agents: an increase in the arrival rate of liquidity providers leads to a substantial increase in local and inside depth as compared to the single market scenario. Thus, our quoted liquidity results show that bid-ask spreads are unambiguously

22

smaller in a single market, whereas the results for top-of-depth depth are ambiguous, such that the overall effect depends on the equilibrium increase in the arrival rate of intermediaries.

Differences in quoted liquidity need not necessarily translate into commensurate differences in transaction costs for traders submitting market orders. Thus, we next compare differences in traded liquidity in the single and fragmented market. We measure traded liquidity by the mean effective spread, which captures the actual transaction costs incurred by traders submitting market orders and is calculated as follows:

$$\text{effective spread} = 2x_t(p_t - m_t)/m_t, \tag{7}$$

where x_t is $+1$ for a buyer-initiated order and -1 for a seller-initiated order, p_t is the transaction price, and m_t is the mid-quote immediately before the transaction. We further decompose effective spread into realized spread and price impact. The former is calculated as follows:

$$\text{realized spread} = 2x_t(p_t - m_{t+k})/m_t, \tag{8}$$

where k is the number of time units in the future. As results are qualitatively similar for k equal to 10, 30 and 60 units, we only report the findings for 30 units. Finally, price impact is defined as:

$$\text{price impact} = \text{effective spread} - \text{realized spread} \tag{9}$$

While the realized spread measures liquidity providers' compensation after accounting for adverse selection losses associated with informed orders, the price impact captures the level of information in a trade. As our model does not contain private information, the price impact captures picking-off risk associated with stale limit orders when new (public) information arrives. Similar to quoted liquidity, we compute inside variants of all three measures of traded liquidity based on inside quote midpoints across the two orders books and local variants based on quote midpoint of the order book where a transaction is executed.

Panel B of Table 3 reports the results. Effective spreads and realized spreads are higher in

a market containing two limit order books. When the arrival rate of zero private value agents remains unchanged, effective spreads decrease from 1.80 ticks and 1.45 ticks based on local and inside quotes, respectively, in a fragmented market, to 1.31 ticks in the single market. The differences are even larger if we double the arrival rate of $\alpha = 0$ agents. Realized inside and local spreads are higher in the fragmented market by approximately 0.15 ticks with the same population of agents, and by 0.5 ticks if we double the participation of intermediaries.

Price impact measured relative to local quotes is lower in a single market whereas it is similar when measured relative to inside quotes. This is because, in fragmented markets, a newly arriving trader is more likely to trade in an order book containing a stale quote, leading to a higher local price impact. Price impacts are lower if we double the arrival rate of $\alpha = 0$ agents, likely because the increased arrival rates increases the probability of being picked off for agents with intrinsic trading motives. The local price impact in this scenario is still slightly larger than that in the single market, though the inside price impact is substantially smaller.

Finally, we analyze the degree of inefficiency in prices. If an asset is traded on two limit order books, the degree of price dislocations may be exacerbated *ceteris paribus*, making prices on individual books less efficient than they would be if all demand and supply were to meet in a single order book. In the context of our model, the effect of these frictions is measured by computing the microstructure noise, defined as the absolute deviation between the quote midpoint m_t and the fundamental value v_t. In Panel C of Table 3, we report mean levels of microstructure noise under the three settings. This value changes from 0.67 ticks in a fragmented market to 0.46 ticks in a single market for local quotes under the baseline parametrization. The corresponding differences based on inside quotes are in the same direction although lower in magnitude. However, in fragmented markets with a higher arrival rate of zero private value agents, the microstructure noise is lower than in the single market, suggesting that prices are more efficient. This is expected as, in the absence of private information, a higher number of traders with no intrinsic reasons to trade results in

a faster adjustment of quotes when public information arrives. In conclusion, similar to the quoted depth result above, the effect depends on the equilibrium increase in the arrival rate of intermediaries in fragmented markets.

4.3 Welfare Analysis

In this subsection, we analyze the economic benefits per agent and for the whole market by examining the effect of fragmentation on welfare. Welfare is measured as the average realized payoff per agent. In addition, we decompose the realized payoffs into gains and losses associated with agents' private values and the trading process.

Suppose that an agent with a private value α enters the market at time t. She submits an order (a limit order or a market order) to any of the books at price $\tilde{p}$ with order direction $\tilde{x}$ (to buy or to sell). Suppose further that the agent does not modify the order and that it is finally executed at time t' when the fundamental value is $v_{t'}$. The agent's realized payoff is then given by:

$$\Pi = e^{-\rho\left(t'-t\right)} \left(\alpha + v_{t'} - \tilde{p}\right) \tilde{x}. \tag{10}$$

We can decompose the agents' payoffs and rewrite equation (10) as:

$$\Pi = Gains\ from\ private\ value + Waiting\ cost\ + Money\ Transfer,\ where$$

$$Gains\ from\ private\ value = \alpha\tilde{x}$$
$$Waiting\ cost = (e^{-\rho(t'-t)} - 1)\alpha\tilde{x} \tag{11}$$
$$Money\ Transfer = e^{-\rho(t'-t)}(v_{t'} - \tilde{p})\tilde{x}$$

The first term in (11), *gains from private value*, represents the gains obtained directly from the intrinsic reasons to trade the asset, $\alpha\tilde{x}$. The second term in Equation (11), *waiting cost*, reflects the cost associated with delaying the realization of *gains from private value*. Agents submitting limit orders do not trade immediately after arriving, but have to wait until the orders' execution. This is costly due to the delay cost ρ. The third term in Equation (11),

25

money transfer, reflects the difference between the fundamental value $v_{t'}$ and the transaction price $\tilde{p}$ discounted back to the arrival time of the agent, and reflects the welfare gain (or loss) associated with the trading process. In general, *money transfer* is related to the immediacy cost incurred when an agent wants to immediately realize her private value. For example, an agent who submits a market order realizes her intrinsic private value without a delay, but she may have to pay a cost for demanding immediacy, which would be reflected in a negative *money transfer*.

Table 4 presents the results. It shows that, under all three parameterizations, the differences in aggregate welfare are negligible (3.745 ticks vs. 3.740 ticks vs. 3.757 ticks), whereas the shifts among categories of agents are substantial.[10] Under the baseline parameterization, intermediaries (agents with $\alpha = 0$) generate lower expected payoffs in a consolidated market than a fragmented market (0.543 ticks vs. 0.626 ticks). When we double their arrival rate, while per agent welfare is lower (0.485 ticks), as a group their welfare is larger ($0.485 \times 2 = 0.970$ ticks) than in the other two scenarios. The difference between a single and fragmented market with constant agent population arises for two reasons. First, there is less price competition in a fragmented market because time priority does not apply across order books. Second, as explained in Section 4.1, intermediaries' expected time to execution in a fragmented market is lower than in a consolidated market, which reduces the delay cost component of their welfare.

Table 4 about here

Agents with $|\alpha| = 4$ and $|\alpha| = 8$ have expected payoffs of 3.510 ticks and 7.265 ticks, respectively, in a single market, whereas the corresponding numbers in a fragmented market are 3.479 and 7.202, respectively. While they do not exhibit any significant difference in their absolute waiting costs under both settings, they incur lower money transfer losses in a

---

[10]In Table 4, under the alternative parametrization, the aggregate per period welfare differs from the average total welfare per trader because traders' per period arrival rate is 1.3 rather than 1.0. In other words, the aggregate per period welfare is $0.485 \times 0.6 + 3.312 \times 0.4 + 7.137 \times 0.3 = 3.757$ ticks.

consolidated market than in a fragmented market. For example, the money transfer losses of agents with private value $|\alpha| = 8$ are -0.572 ticks in a consolidated market versus -0.626 ticks in a fragmented market. This is because lower price competition in a fragmented setting leads to higher immediacy costs for the market orders they submit.[11] When we double the arrival rate of agents with $\alpha = 0$, agents with $\alpha \neq 0$ change their behavior by relying more on market orders (see Table 1). This reduces their overall execution time, resulting in lower waiting costs, but the switch to market orders results in higher money transfer losses (for example, -0.835 ticks for agents with $|\alpha| = 8$).

In conclusion, as aggregate welfare effects in the model are small, any interpretation regarding the overall desirability of fragmentation needs to go beyond the model. While market participation in the model is exogenous, in real markets, one would expect that agents endogenously enter the market as long as expected profits from participation are higher than the entry costs. Higher intermediation profits in fragmented markets should lead to increased participation by liquidity providers. If market participation by intermediaries is costly – a realistic assumption considering expenses such as colocation charges, subscriptions to data feeds, investments in high speed networks – these agents incur costs while not contributing to an increase in aggregate welfare. In other words, their entry decision, while privately optimal, is socially suboptimal.

5. Empirical Application

In this section, we empirically analyze some of the predictions generated by our model. To this end, we conduct an event study based on Euronext's implementation of a single order book per asset for its Paris, Amsterdam, and Brussels markets.

---

[11]Note that in Table 4, the total money transfers do not add up to zero, as the expected payoff in a single transaction of the limit order and the corresponding market order are discounted back to different times. This is due to the asynchronous arrivals of the agents who submit these two orders. However, the instantaneous money transfer not discounted back is equal to zero.

5.1 Euronext's Institutional Background

Euronext was formed in 2000 following a merger of the Paris, Amsterdam and Brussels stock exchanges. In 2002, the Lisbon Stock Exchange became part of Euronext.[12,13] Stock listings on Euronext pertain to a listing on one or more national markets.[14] Until 13 January 2009, each national listing corresponded to the operation of an independent limit order book. For example, a stock listed on the Paris market would be traded on the limit order book of Euronext Paris. Firms cross-listed in multiple Euronext markets traded on multiple Euronext order books in parallel, besides other competing markets.[15]

On 16 August 2007, the exchange announced its intention to eliminate this arrangement for their Paris, Amsterdam and Brussels markets by consolidating all trading in these markets in a single order book, the so-called "Market of Reference" (MoR). Cross-listed firms had to choose one MoR that continued operating after the implementation of the single order book. This new arrangement was implemented on 14 January 2009.

The existence of multiple order books led to fragmentation of order flow routed to Euronext. As the rules and trading protocols governing the individual order books were identical, the introduction of a single order book decreased fragmentation of trading in these cross-listed stocks without any corresponding change in the competitive environment.[16] This is particularly relevant as it allows us to test the model predictions directly. Euronext, in its press release announcing the event, made it clear that, for investors, the trading environment remained unchanged: "The Single Order Book will have no impact on the NSC system as

---

[12]Nielsson (2009) studies the effects on the formation of Euronext and the later addition of the Lisbon Stock Exchange on the liquidity of listed firms and finds positive effects on the liquidity of large firms.

[13]In 2007, Euronext merged with the NYSE to form NYSE Euronext, which was taken over by Intercontinental Exchange in 2012. In 2014, Euronext was spun off through an IPO.

[14]With the implementation of the Markets in Financial Instruments Directive (MiFID), all rules prohibiting trading outside the national markets were repealed such that investors can now trade these stocks in any market.

[15]Lescourret and Moinas (2018) exploit these cross-listings within Euronext to study market-making in fragmented markets.

[16]Pagano and Padilla (2005) highlight the steps taken by Euronext to standardize its trading protocols and technological platform. They further argue that, the efficiency gains generated through the merger were a direct consequence of these steps.

the market rules and order book management will remain unchanged [. . .] In practice, from a trading perspective, Single Order Book implementation simply means the end of order book trading on marketplaces other than the market of reference."[17] Moreover, as market consolidation was based on a business decision by Euronext, all multi-listed stocks received the same treatment such that there was no selection bias. Finally, the announcement was made more than one year before the event date in order to allow market participants to adapt and test their trading systems. This eliminates potential concerns about the event date confounding with other market events around the same time.[18] Thus, the empirical setting of the transition from a multi-market environment to a single market can be viewed as a quasi-natural experiment, allowing us to test the hypotheses obtained from our theoretical model.

5.2   Sample Selection

A total of 45 instruments, cross-listed on at least two of the three Euronext markets, were affected by the event. However, we reduce the sample of treated stocks used in our study for several reasons. First, we remove stocks whose primary listing is not on Euronext. These include stocks whose main trading activity takes place in other European markets or in the United States. Second, we eliminate exchange-traded mutual funds because we do not expect their trading activity to be comparable to that of listed firms. Finally, we require that the total share of trading activity on the less active Euronext order book be at least equal to 1% of the stock's total Euronext trading volume. This reduces the list of instruments to ten. We further exclude one additional stock because of data errors, reducing our final sample to nine stocks.[19] The number of stocks is small due to the unique nature of the event we

---

[17]See Euronext press release dated 14 January 2009 available here.

[18]Although the original date of implementation was postponed, this was due to technical reasons as opposed to concerns about market conditions. The final implementation date was announced more than 60 days in advance.

[19]One stock in our sample was listed in all three Euronext markets. However, we exclude the least active limit order book as it had a market share of 0.3%.

study. Nonetheless, our sample consists of the whole population of stocks affected by the event, except a subset of stocks which are excluded through objective criteria.

As suggested by Davies and Kim (2009), we construct a matched control group of stocks based on stock price and market capitalization obtained from Compustat Global using the distance metric employed by Huang and Stoll (1996). Specifically, for each stock in our treatment group, we identify the stock that is its closest match in terms of these two criteria as on the last trading day of 2008 (30 December 2008). The population of stocks from which the control group is constructed comprises all stocks with a primary listing on a single Euronext market. In other words, these firms are not affected by the event.[20]

Using a control group allows us to identify the effects of reduced fragmentation, implicitly controlling for market-wide changes in variables such as liquidity and volatility. It also allows us to control for two additional market-wide changes implemented by Euronext close in time to the introduction of a single order book. First, a harmonized settlement platform known as the Euroclear Settlement for Euronext-zone Securities for all French, Dutch and Belgian stocks was implemented on 19 January 2009. Second, the Universal Trading Platform, having "superior functionality, faster speed and much greater capacity", was introduced on 16 February 2009.[21] These were market-wide events that affected both the control and treatment stocks. Consequently, we can attribute any differences in trading activity and market quality between the two groups exclusively to market consolidation resulting from the introduction of a Single Order Book.

For the purpose of our analyses, we define all days from 1 December 2008 to 13 January 2009 as the pre-event period and all days from 19 January 2009 to 27 February 2009 as the

---

[20]As our sample period coincides with the global financial crisis, in an alternative matching procedure we additionally constrain the choice of control stocks to financial (non-financial) firms for financial (non-financial) treatment firms. Euronext also implemented price contingent tick sizes for the largest firms listed in its four national markets on 15 December 2008. As three firms in the control group were affected by this change, we test our results using alternative control firms that were unaffected by this event. The results of these untabulated robustness tests are similar to those reported in this section.

[21]See press release titled "NYSE Euronext's European Equities Trading Successfully Migrates to the Universal Trading Platform" dated 17 February 2009 available here.

post-event period. We exclude all trading days from the event date until the end of the week in order to eliminate any effect associated with the transition. We also exclude 24 December 2008 and 31 December 2008 from the pre-event period.

5.2.1 Data Description and Summary Statistics

We use high-frequency trades and best quotes data from Thomson Reuters Tick History between December 2008 and February 2009 for the purpose of our analysis. The data is time-stamped with a millisecond resolution. As our focus is on the continuous trading session, we exclude data before 09:01 and after 17:29 on each day. This means that we exclude trades during the opening and closing auctions, and within the first and last minute of the continuous trading session.

Table 5 describes the characteristics of stocks in the treatment and control group. The average market capitalization across stocks in the treatment (control) group is €4.4 (€4.8) billion and the average stock price is €18.4 (€18.2). This highlights the high matching quality between firms in the treatment and control groups. The share of the more active trading venue as a percentage of total Euronext volume across all the days in the pre-event period ranges from 54% to 98% across the nine stocks in the treatment group. The simple (volume-weighted) average across all stocks is 78% (62%). This implies that substantial volume was executed on the less active market. The market share of the sole listing Euronext venue for the stocks in the control group is, by construction, 100%. Trading activity when measured in terms of the number of trades (not reported) provides a similar picture.

Table 5 about here

5.2.2 Estimation Methodology

In order to test the main implications of our model, we compute several variables capturing trading activity, liquidity and price efficiency. Similar to our theoretical results, we compute

31

local and inside measures. Unsophisticated investors who choose to trade only on a single order book are likely to select the more active and liquid book. Hence, we compute the local measures for the market having a higher trading volume during the pre-event period. We estimate a panel difference-in-differences regression with stock and day fixed effects and standard errors double-clustered by stock and day. We estimate the regression for levels and natural logarithms of the variables of interest in order to account for the wide dispersion in levels across the stocks in our sample.

5.3 Empirical Results

5.3.1 Quoted Liquidity

We begin by analyzing the effect on time-weighted quoted spread and top-of-book depth. Table 6 presents the results. Consistent with our theoretical findings, we observe an overall improvement in quoted spreads on Euronext after the introduction of a single order book. The more active of the two Euronext markets experiences a significant reduction in local quoted spreads of 160bps or approximately 35%. The effect on inside spreads depends on the test specification and is statistically insignificant. The absence of a significant improvement in the inside spread can be explained by the fact that in real markets, different from our model, market participants do not always route their orders optimally, i.e., to the market offering the best price.[22] Thus, while in the model inside spreads correctly reflect the gains, before adverse selection, that liquidity providers expect to earn, an empirically observed non-zero probability of traders routing their orders sub-optimally allows liquidity providers to earn larger trading profits. This effect disappears after consolidation, *ceteris paribus* leading to an increase in inside quoted spreads. Conversely, an increase in price competition among liquidity providers, as predicted by our theory, leads to a decrease in quoted spreads after

---

[22]In European markets, best execution requirements allow brokers to consider other criteria besides price when making order routing decisions. In contrast to the US, European markets also do not have an order protection rule that requires exchanges to re-route orders to venues offering a superior price. Even in the US, communication latencies between geographically dispersed exchanges and exceptions to the order protection rule result in liquidity takers sometimes obtaining sub-optimal prices.

consolidation. These two effects cancel out leading to insignificant coefficients for inside spreads.

<p align="center">*Table 6 about here*</p>

We further observe a positive though statistically insignificant effect of order flow consolidation on local and inside top-of-book depth. These results lie in between those observed in the simulations with different participation rates of intermediaries. In other words, they are consistent with the amount of market-making in fragmented markets being larger than, but less than twice as large as in a consolidated market.

5.3.2 Traded Liquidity

Table 7 presents the results for trade size-weighted effective spread and its decomposition into realized spread and price impact based on 30 and 60 second future mid-quotes. Effective inside (local) spreads decrease by an economically large 34bps (76bps) after the introduction of a single order book and realized log spreads decrease by anywhere between 47% and 68% depending on the time horizon and whether we use local or inside quotes. The decrease in log spreads is always significant. These results are consistent with our theoretical predictions.

<p align="center">*Table 7 about here*</p>

The observed change in inside price impacts differs depending on the empirical specification and is never significant, whereas our theory tells us the effect should be near-zero or positive. The fact that liquidity takers empirically sometimes do not trade on the market offering the best price may explain why inside price impacts are not larger in the fragmented market. An order trading against a standing limit order at a price inferior to the lowest ask or highest bid does not mechanically generate an inside price impact even if it executes against the entire limit order, leading to a relatively smaller inside price impact compared to

<p align="center">33</p>

the theoretical prediction. Additionally, the empirical results, in contrast to the model, also capture the effects associated with private information possessed by traders. Lower transaction costs potentially allow traders with small amounts of private information to profitably participate, leading to a decrease in average price impact. The latter channel may cancel out the positive effect of consolidation on price impact predicted by our theory.

5.3.3 Price Efficiency

In our simulations, we examine the price efficiency by measuring the extent to which the mid quote deviates from the fundamental value. Empirically, as we cannot observe the fundamental value, we measure price efficiency using return autocorrelations and variance ratios. Return autocorrelations are measured at 30 second and 5 minute intervals. Variance ratios capture the deviation between long-term and short-term return variance and are calculated as one minus the ratio of long-term and short-term return variance, each scaled by the respective time periods. We calculate variance ratios between 30 second and 5 minute return variances. As in Boehmer and Kelley (2009), we compute the effect on absolute values of both measures because we are interested in departures from a random walk in either direction. The closer these measures are to zero, the more closely the price path resemble a random walk.

Table 8 presents the results on price efficiency.[23] The variance ratio decreases after the implementation of the single order book and the change is statistically significant. The results for autocorrelations also point to improved price efficiency in a single market compared to a fragmented market, but the change is insignificant. These results generally provide evidence for unchanged or higher price efficiency in consolidated markets, which is consistent with our theoretical predictions. When compared to the latter, the empirical evidence appears consistent with a fragmented market containing more, but less than twice as much

---

[23]The empirical results based on returns measured at other frequencies are qualitative similar to those reported here.

intermediation as a consolidated market.

<center>*Table 8 about here*</center>

5.3.4 Trading Volume and Arbitrage

The existence of multiple order books allows market participants to earn arbitrage profits by exploiting instances of crossed markets, i.e., situations where the bid price on one order book is higher than the ask price on the other. Whether such instances are resolved through a quote adjustment or a trade does not make a difference for price efficiency. However, the latter leads to losses for limit order traders who, in turn, increase their bid-ask spreads. This arbitrage-driven rent extraction may lead to welfare losses if otherwise beneficial trades are crowded out (Foucault et al., 2017; Budish et al., 2015).

We measure trades potentially associated with arbitrage activity. We start by identifying instances of a crossed order book. Such a situation can arise as a result of new order(s) submitted to either or both order book(s). We then identify whether these instances are resolved through a trade, quote-update, or both. This approach is similar to Foucault et al. (2017) who define the resolution through trades as toxic arbitrage if the following two conditions are fulfilled: (i) prices offered in different markets allow aggressive traders to earn a profit by trading against the bid on one market and ask on the other; (ii) they are able to do so because of liquidity providers' slow reaction to new information, rather than them offering attractive prices to manage their inventories. Fragmentation is an obvious precondition for such arbitrage trades to occur. For each stock-day, we calculate the number of unique crossed instances, the fraction of the day when inside spreads are negative, and the total trading volume contributing to the resolution of a crossed market. Panel A of Table 9 reports the mean values for each stock across all days in the pre-event period. The frequency of unique instances of a crossed market for an average day ranges from 0.3 to 717 across all stocks, with an average value of 141, which corresponds to one instance every 3.6

<center>35</center>

minutes. An average stock has a negative inside spread for 5.6% of the continuous trading session. Finally, 6.2% of the total trading volume on Euronext for an average stock can be attributed to the resolution of instances where the two markets are crossed. Approximately 50% of this, or 3% of total Euronext trading volume, is associated with toxic arbitrage as defined in Foucault et al. (2017).[24]

Table 9 about here

As arbitrage trades between multiple Euronext order books are mechanically eliminated after the introduction of a single order book, trading volume should, everything else equal, be reduced. Panel B of Table 9 indicates that the effect of introducing a single order book is in fact slightly positive but insignificant, suggesting that the volume transacted by investors with intrinsic motives to trade increases after consolidation. This indicates that such investors who were earlier crowded out in a less liquid fragmented market now gainfully participate, leading to an overall welfare gain after consolidation. This is consistent with our theory. In the model, the volume traded by agents with intrinsic reasons to trade remains constant because their private values are assumed to be sufficiently large such that they never refrain from trading.

5.3.5 Testing for Parallel Trends

In a difference-in-differences regression, the existence of parallel trends between the treatment and control groups before imposition of the treatment is a key identifying assumption. We employ two tests to validate the existence of parallel trends. We first plot the time trends of the different variables to visually identify any obvious violations of parallel trends. Second, as suggested by Pischke (2005), we augment our original difference-in-dfferences regression by including the interactions between the treatment variable and six weekly dummies on

---

[24]As timestamps in TRTH are not exchange timestamps, we rerun our analysis by ignoring arbitrage opportunities that are resolved in less than 100 milliseconds. While the frequency of arbitrage opportunities and the corresponding trading volume is obviously smaller, our results remain qualitatively unchanged.

each side of the event date.[25] If the parallel trends assumption is satisfied, then all the weekly pre-event dummies should be jointly insignificant. In Figure 2 and Table 10, we report the plots and regression results for effective spreads and realized spreads as our key theoretical predictions concern these two variables and our empirical analysis also produces the strongest results for these. Based on the above tests, we conclude that the assumption of parallel trends in the pre-event window holds.[26]

Figure 2 about here

Table 10 about here

6. Conclusion

We examine the effects of market fragmentation when competition between trading venues is non-existent. Such fragmentation is routinely observed after exchange mergers, when a single exchange operator continues operating multiple order books to trade the same asset post merger. In an attempt to extract synergies from the merger, the operator typically eliminates most structural and technological differences across the markets resulting in operator-level order flow fragmenting across (nearly) identical limit order books.

We model a fragmented market consisting of two identically-organized limit order books populated by heterogeneous agents who endogenously choose to supply or consume liquidity. Our model allows us to examine the effects on several aspects of market performance such as liquidity, price efficiency, agents' payoffs and overall welfare. As time priority is not enforced across markets, fragmentation leads to reduced price competition between intermediaries. This results in the deterioration of liquidity in fragmented markets as compared to the

---

[25]The weekly dummies W_k are defined based on calendar weeks with $k \in \{-6, -5, -4, \ldots, 4, 5, 6\}$. The week immediately before the event date W_{-1} includes three additional days before the event date.

[26]It is worth noting that our sample period coincides with a market-wide deterioration in liquidity during the depth of the global financial crisis. This is visible in Figure 2 as an increase in effective and realized spreads for the control firms.

consolidated market benchmark. While overall welfare is nearly identical under both market setups, the distribution of welfare across the heterogeneous agent types in the model is markedly different. Agents with intrinsic trading motives extract lower payoffs in fragmented markets whereas agents acting as intermediaries are better off in fragmented markets.

These higher intermediation gains should, under conditions of endogenous entry, lead to more intermediaries entering the market. We mimic these conditions by doubling the population of intermediaries in the model while keeping all other market parameters constant. We observe that under these conditions the allocation of trading gains between intermediaries and non-intermediaries shifts further in favor of the former without materially altering overall welfare. These results suggest that fragmentation leads to investment in intermediation capacities, such as the high-speed connections required to access the trading systems and real-time data feeds from exchanges, that is socially wasteful.[27]

We empirically test the model implications by investigating the effects of Euronext's decision to introduce a single order book for their Paris, Amsterdam, and Brussels markets. As opposed to the existing empirical literature which necessarily investigates the joint impact of changes in fragmentation and competition (say, when a new operator enters the market), this event allows us to examine the effects associated with the consolidation of multiple non-competing order books. The empirical analysis broadly confirms the theoretical predictions related to the effects on liquidity, price efficiency, and limit order traders' profits. Additionally, we also obtain evidence that trading volume after consolidation does not decrease even though cross-market arbitrage opportunities are eliminated. This suggests that, while the revenues generated by modern exchanges' from the sale of market data may decrease after consolidation, improvements in market quality need not come at the expense of reduced trading fees for the exchange operators.

Overall our results suggest that the positive externalities associated with consolidating

---

[27]Cespa and Foucault (2013) and Easley et al. (2016) further highlight the adverse effects associated with exchanges providing differential access to market data feeds.

order flow in a single location (or fewer locations) still exist and are substantial. This is true even in modern electronic limit order markets where the activities of high-frequency traders serve to integrate fragmented order books. The adverse effects of fragmentation are significantly larger for unsophisticated investors who do not possess the technological ability to route their trades to the most advantageous order book. For such investors consolidation of order flow, at least between non-competing markets, likely results in transaction cost reductions.

Our results also have important policy implications. Regulators may be able to improve the welfare of investors who trade for intrinsic motives by: (i) preventing individual market operators from keeping an artificially high(er) level of order flow fragmentation in the absence of commensurate benefits; and (ii) limiting excessive investment in intermediation capacities necessary to link multiple order books which come at a cost to end investors.

39

References

Amihud, Y., B. Lauterbach, and H. Mendelson (2003). The value of trading consolidation: evidence from the exercise of warrants. *Journal of Financial and Quantitative Analysis 38*(04), 829–846.

Baldauf, M. and J. Mollner (2018). Trading in Fragmented Markets. *Working Paper*.

Boehmer, B. and E. Boehmer (2003). Trading your neighbor's ETFs: Competition or fragmentation? *Journal of Banking and Finance 27*(9), 1667–1703.

Boehmer, E. and E. K. Kelley (2009). Institutional Investors and the Informational Efficiency of Prices. *Review of Financial Studies 22*(9), 3563–3594.

Budish, E., P. Cramton, and J. Shim (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal of Economics 130*(4), 1547–1621.

Cespa, G. and T. Foucault (2013). Sale of Price Information by Exchanges: Does It Promote Price Discovery? *Management Science 60*(1), 148–165.

Chao, Y., C. Yao, M. Ye, and R. Greenwood (2018). Why Discrete Price Fragments U.S. Stock Exchanges and Disperses Their Fee Structures. *Review of Financial Studies*.

Chlistalla, M. and M. Lutat (2011). Competition in securities markets: the impact on liquidity. *Financial Markets and Portfolio Management 25*(2), 149–172.

Chowdhry, B. and V. Nanda (1991). Multimarket trading and market liquidity. *Review of Financial Studies 4*(3), 483–511.

Chung, K. H. and C. Chuwonganant (2012). Regulation NMS and Market Quality. *Financial Management 41*(2), 285–317.

Colliard, J.-E. and T. Foucault (2012). Trading Fees and Efficiency in Limit Order Markets. *Review of Financial Studies 25*(11), 3389–3421.

Davies, R. J. and S. S. Kim (2009). Using matched samples to test for differences in trade execution costs. *Journal of Financial Markets 12*(2), 173–202.

Degryse, H., F. de Jong, and V. v. Kervel (2015). The Impact of Dark Trading and Visible Fragmentation on Market Quality. *Review of Finance 19*(4), 1587–1622.

Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in io. *Handbook of Industrial Organization 3*, 1887–1966.

Easley, D., M. O'Hara, and L. Yang (2016). Differential Access to Price Information in Financial Markets. *Journal of Financial and Quantitative Analysis 51*(4), 1071–1110.

Foucault, T., O. Kadan, and E. Kandel (2013). Liquidity cycles and make/take fees in electronic markets. *Journal of Finance 68*(1), 299–341.

Foucault, T., R. Kozhan, and W. W. Tham (2017). Toxic Arbitrage. *Review of Financial Studies 30*(4), 1053–1094.

Foucault, T. and A. J. Menkveld (2008). Competition for Order Flow and Smart Order Routing Systems. *Journal of Finance 63*(1), 119–158.

Glosten, L. R. (1998). Competition, Design of Exchanges And Welfare. *Working Paper*.

Goettler, R. L., C. A. Parlour, and U. Rajan (2005). Equilibrium in a Dynamic Limit Order Market. *Journal of Finance 60*(5), 2149–2192.

Goettler, R. L., C. A. Parlour, and U. Rajan (2009). Informed traders and limit order markets. *Journal of Financial Economics 93*(1), 67–87.

Gomber, P., S. Sagade, E. Theissen, M. C. Weber, and C. Westheide (2017). Competition between equity markets: A review of the consolidation versus fragmentation debate. *Journal of Economic Surveys 31*(3), 792–814.

Gresse, C. Effects of Lit and Dark Market Fragmentation on Liquidity. *Journal of Financial Markets 35*, 1.

Harris, L. E. (1993). Consolidation, Fragmentation, Segmentation and Regulation. *Financial Markets, Institutions & Instruments 2*(5), 1–28.

Haslag, P. H. and M. Ringgenberg (2015). The Causal Impact of Market Fragmentation on Liquidity.

Hengelbrock, J. and E. Theissen (2009). Fourteen at One Blow: The Market Entry of Turquoise. *Working Paper*.

Huang, R. D. and H. R. Stoll (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics 41*(3), 313–357.

Lescourret, L. and S. Moinas (2018). Fragmentation and Strategic Market-Making. *Working Paper*.

Madhavan, A. (1995). Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies 8*(3), 579–603.

Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory 100*(2), 191–219.

Mendelson, H. (1987). Consolidation, Fragmentation, and Market Performance. *Journal of Financial and Quantitative Analysis 22*(02), 189–207.

Nguyen, V., B. F. Van Ness, and R. A. Van Ness (2007). Short- and Long-Term Effects of Multimarket Trading. *Financial Review 42*(3), 349–372.

Nielsson, U. (2009). Stock exchange merger and liquidity: The case of euronext. *Journal of Financial Markets 12*(2), 229–267.

Pagano, M. (1989). Trading Volume and Asset Liquidity. *Quarterly Journal of Economics 104*(2), 255–274.

Pagano, M. and J. A. Padilla (2005). Gains from Stock Exchange Integration: The Euronext Evidence. *Working Paper*.

41

Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse'of dimensionality. *Econometrica 69*(5), 1261–1281.

Panayides, M. A., B. Rindi, and I. M. Werner (2017). Trading Fees and Intermarket Competition. *Working Paper*.

Parlour, C. A. and D. J. Seppi (2003). Liquidity-Based Competition for Order Flow. *Review of Financial Studies 16*(2), 301–343.

Pischke, J.-S. (2005). Empirical Methods in Applied Economics Lecture Notes.

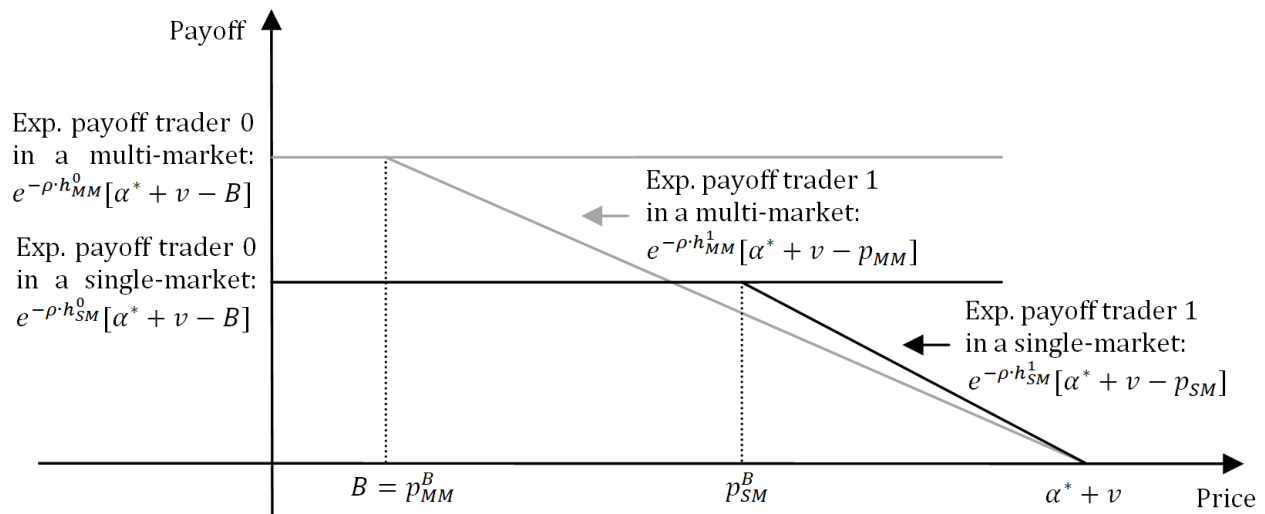Rieder, U. (1979). Equilibrium plans for non-zero-sum markov games. *Game Theory and Related Topics*, 91–101.

Figure 1. Partial Equilibrium with Two Competing Traders

This figure shows the partial equilibrium with two traders under competition in a single-market (black lines) and in a fragmented market (gray lines). Trader 1 wants to submit a limit buy order and observes that other trader, labeled Trader 0, previously submitted a limit order at price B. The horizontal lines represent Trader 0's expected payoff in the case that she keeps her order at price B after the order submission of Trader 1. The sloped lines reflect the expected payoff of Trader 1 from the execution of her limit buy order. The slopes of the expected payoff lines of Trader 1 are $e^{-\rho \cdot h_{SM}^1}$ and $e^{-\rho \cdot h_{MM}^1}$ for the single and fragmented market setup, respectively.

43
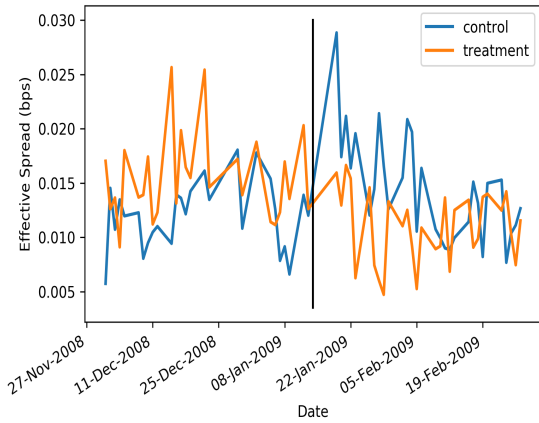
(a) Effective Local Spread
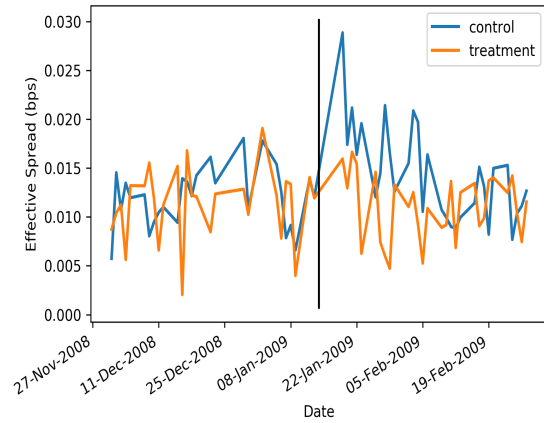
(b) Effective Inside Spread
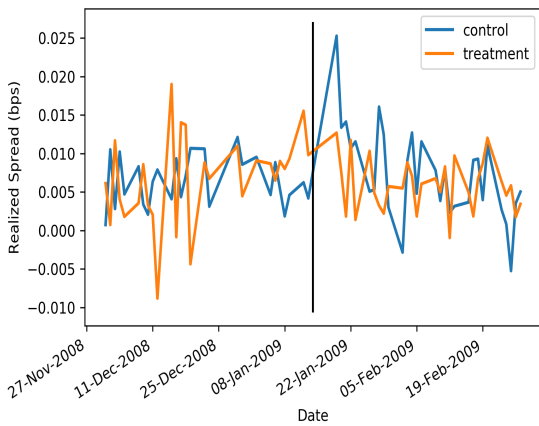
(c) Realized Local Spread
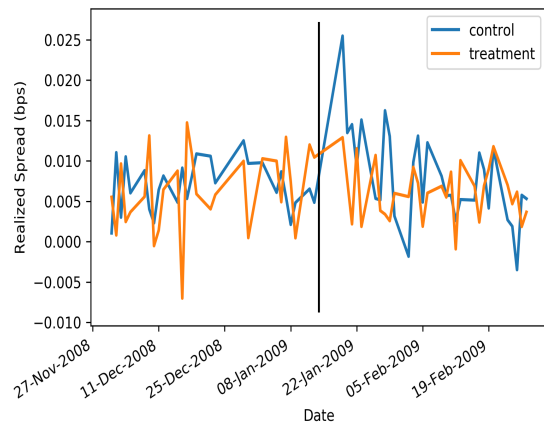
(d) Realized Inside Spread

Figure 2. Testing for Parallel Trends

We plot the local and inside versions of daily average effective and realized spreads between 1 December 2008 and 27 February 2009 for the treatment and control firms. The vertical line coincides with 14 January 2009, the day Euronext implemented a single order book.

44

Table 1. Trader Behaviour: Choice of Order Type

This table reports (1) the probability of observing a limit order at the bid or ask prices, calculated as the number of limit orders placed at the bid or ask prices over the total number of limit orders in the book, and (2) the distribution of limit and market orders executed by each agent type. We report these frequencies for a single and fragmented market under the baseline parameterization and for a fragmented market with the population of agents with $\alpha = 0$ doubled.

Market Setup	Order	Best price	Private value $|\alpha|$		
	Type	Prob.	0	4	8
Single Market	Limit	50.0%	77.1%	52.5%	19.6%
	Market	–	22.9%	47.5%	80.4%
Frag Market	Limit	41.6%	78.4%	50.7%	20.7%
	Market	–	21.6%	49.3%	79.3%
Frag Market Double $\alpha = 0$	Limit	36.0%	90.0%	23.4%	5.6%
	Market	–	10.0%	76.6%	94.4%

45

Table 2. Trader Behaviour: Limit Order Statistics

This table reports statistics concerning the limit orders submitted by each agent type. Panel A reports the time to execution, which is defined as the time when an order executes minus the market entry time of the agent submitting the limit order. Panel B reports the limit price aggressiveness of executed orders i.e., the difference between the limit price and the asset's fundamental value at the time of submission (in ticks). Panel C reports the picking-off risk i.e., the frequency of limit orders executed with a negative instantaneous payoff. We report all statistics for a single and fragmented market under the baseline parameterization and for a fragmented market with the population of agents with $\alpha = 0$ doubled.

Market Setup	Private value $|\alpha|$			Total
	0	4	8	
Panel A: Time to Execution				
Single Market	14.9	3.5	2.1	8.6
Frag Market	11.5	3.7	1.9	7.1
Frag Market Double $\alpha = 0$	15.4	2.1	1.0	13.1
Panel B: Limit Price Setting Behavior				
Single Market	1.03	-0.10	-0.86	0.34
Frag Market	1.06	-0.08	-0.88	0.35
Frag Market Double $\alpha = 0$	1.05	-0.06	-1.13	0.83
Panel C: Picking-Off Risk				
Single Market	4.1%	26.8%	73.9%	21.8%
Frag Market	3.0%	25.8%	72.1%	20.8%
Frag Market Double $\alpha = 0$	6.0%	26.9%	80.1%	10.9%

46

Table 3. Impact on Market Quality

This table compares the impact of fragmentation on the quality of markets. Panels A, B, and C report the differences in quoted liquidity, traded liquidity, and price efficiency, measured by microstructure noise. We compute local and inside versions of all measures. The former is based on the local quotes, i.e., the bid and the ask prices of a local market, whereas the latter is based on inside quotes, i.e., the highest bid and the lowest ask across the two limit order books. In the single market setting the local and inside quotes are identical. Quoted spread is the difference between the best bid and best ask prices. Quoted depth is the average total number of limit orders waiting to be executed at the top of the order book. Effective spread is defined in Equation (7) and realized spread is defined in Equation (8). Price impact is the difference between effective spread and realized spread. Microstructure noise is the absolute difference between the quote midpoint and the fundamental value v_t. We report all statistics for a single and fragmented market under the baseline parameterization and for a fragmented market with the population of agents with $\alpha = 0$ doubled. We omit standard errors for the differences because a large number of trader arrivals leads to difference in means to the order of 10^{-2} to be statistically significant.

	Single Market	Frag Market	Frag Market Double $\alpha = 0$	Difference	
	(1)	(2)	(3)	(1) − (2)	(1) − (3)
Panel A: Quoted Liquidity					
Quoted Spread: Local	1.565	2.601	2.240	−1.035	−0.675
Quoted Spread: Inside	1.565	1.904	1.860	−0.339	−0.295
Quoted Depth: Local	1.584	1.082	1.692	0.502	−0.107
Quoted Depth: Inside	1.584	1.445	2.751	0.139	−1.166
Panel B: Traded Liquidity					
Effective Spread: Local	1.312	1.799	1.862	−0.487	−0.550
Effective Spread: Inside	1.312	1.452	1.613	−0.140	−0.301
Realized Spread 30: Local	0.865	1.013	1.372	−0.148	−0.506
Realized Spread 30: Inside	0.865	1.011	1.371	−0.146	−0.506
Price Impact 30: Local	0.441	0.789	0.487	−0.348	−0.046
Price Impact 30: Inside	0.441	0.442	0.242	−0.001	0.199
Panel C: Microstructure Noise ($|v_t - m_t|$)					
Local	0.464	0.670	0.369	−0.207	0.095
Inside	0.464	0.570	0.350	−0.106	0.114

Table 4. Decomposition of Welfare by Trader Type

This table reports the welfare defined as the average realized payoff. In addition, this table presents waiting cost and money transfer defined in Equation 11. These measures are reported for each agent type. We report all statistics for a single and fragmented market under the baseline parameterization and for a fragmented market with the population of agents with $\alpha = 0$ doubled. The last column reports the average aggregate welfare per period, which differs from the average welfare per trader under the alternative parameterization, because the arrival rate of traders per period is 1.3 rather than 1.0. All measures are reported in ticks.

	Average welfare per trader				Waiting cost per trader				Money transfer per trader				Aggr. Welfare per period
	Abs. Private Value			Total	Abs. Private Value			Total	Abs. Private Value			Total	
	0	4	8		0	4	8		0	4	8		
Single Market	0.543	3.510	7.265	3.745	0.000	-0.350	-0.162	-0.189	0.543	-0.140	-0.572	-0.065	3.745
Frag Market	0.626	3.479	7.202	3.740	0.000	-0.355	-0.172	-0.193	0.626	-0.166	-0.626	-0.066	3.740
Frag Market $2 \times \alpha = 0$	0.485	3.312	7.137	2.890	0.000	-0.127	-0.029	-0.046	0.485	-0.561	-0.835	-0.142	3.757

48

Table 5. Stock Characteristics

This table reports the characteristics of the treatment stocks and the corresponding control group of stocks. Market Capitalization is the product of shares outstanding and Stock Price as on 31 December 2008. Volume % Large is the share in total Euronext trading volume, between 1 December 2008 and 13 January 2009, of the most active limit order book on Euronext.

	Treatment Stocks				Control Stocks		
Stock	Market Cap € million	Stock Price €	Volume % Large	Stock	Market Cap € million	Stock Price €	Volume % Large
DEXI	3,355	2.9	54.3%	STM	3,355	4.6	100.0%
FOR	2,187	0.9	71.5%	CNAT	3,635	1.3	100.0%
ISPA	24,985	17.2	55.6%	ABI	25,439	15.9	100.0%
UNBP	8,598	104.9	87.7%	HRMS	10,652	101	100.0%
GLPG	80	3.8	77.1%	OMT	84	4	100.0%
ONCOB	87	6.6	90.7%	TAM	81	6.8	100.0%
RCUS	193	6.2	98.0%	AMG	184	6.9	100.0%
VRKP	105	20	94.4%	SMTPC	117	20	100.0%
THEB	59	3.5	68.4%	DEVG	63	3.5	100.0%
MEAN	4,405	18.4	77.5%		4,846	18.2	100.0%

49

Table 6. Empirical Findings: Impact on Quoted Liquidity

This table presents the impact of introducing a single order book on time-weighted quoted bid-ask spreads and quoted top-of-book depth. We report the effects on both local and inside liquidity. The former are based on the bid and ask prices of the order book with a larger trading volume in the pre-event period, whereas the latter are based on inside quotes, i.e., the highest bid and the lowest ask across the two limit order books. We estimate a difference-in-difference regression for the level and logarithm of the dependent variable, and report the coefficient of the interaction term which equals one for all treatment stocks and for all days on or after 19 January 2009, and zero otherwise. We employ stock and day fixed effects and double cluster standard errors by stock and day. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

| | Treatment | Control | Effect Size | |
	Post-Pre	Post-Pre	Levels	Logs
Quoted Spread: Local	−0.013	0.281	−0.016*	−0.351***
Quoted Spread: Inside	0.011	0.281	−0.272	0.078
Quoted Depth: Local	703	−3,298	4,009	0.175
Quoted Depth: Inside	278	−3,298	3,585	0.068

50

Table 7. Empirical Findings: Impact on Traded Liquidity

This table presents the impact of introducing a single order book on the trade size-weighted effective spreads, realized spreads, and price impact. We report the effects on both local and inside liquidity. The former are based on the quotes in the order book where the trade is executed, whereas the latter are based on inside quotes, i.e., the highest bid and the lowest ask across the two limit order books. Realized spreads and price impact are measured based on mid-quotes 30 and 60 seconds after a trade. We estimate a difference-in-difference regression for the level and logarithm of the dependent variable, and report the coefficient of the interaction term which equals one for all treatment stocks and for all days on or after 19 January 2009, and zero otherwise. We employ stock and day fixed effects and double cluster standard errors by stock and day. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Effective Spread: Local	−0.448	0.240	−0.759**	−0.267**
Effective Spread: Inside	−0.035	0.240	−0.341*	−0.054
Realized Spread 30: Local	−0.087	0.104	−0.231	−0.470**
Realized Spread 60: Local	−0.116	0.070	−0.226	−0.682**
Realized Spread 30: Inside	−0.036	0.104	−0.184	−0.494**
Realized Spread 60: Inside	−0.049	0.070	−0.161	−0.657**
Price Impact 30: Local	−0.361	0.136	−0.529	−0.161
Price Impact 60: Local	−0.332	0.171	−0.533	−0.171
Price Impact 30: Inside	0.001	0.136	−0.158	0.153
Price Impact 60: Inside	0.014	0.171	−0.181	0.108

Table 8. Empirical Findings: Impact on Price Efficiency

This table presents the impact of introducing a single order book on return autocorrelation and variance ratios, which are used as proxies for prices efficiency. We report the effects on both local and inside price efficiency. The former are based on the bid and ask prices of the order book with a larger trading volume in the pre-event period, whereas the latter are based on inside quotes, i.e., the highest bid and the lowest ask across the two limit order books. Autocorrelations are calculated with returns measured at 30-second and 5-minute intervals and variance ratios are based on 30-second and 5-minute returns. We estimate a difference-in-difference regression for the level and logarithm of the dependent variable, and report the coefficient of the interaction term which equals one for all treatment stocks and for all days on or after 19 January 2009, and zero otherwise. We employ stock and day fixed effects and double cluster standard errors by stock and day. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Autocorrelation 30: Local	−0.014	0.006	−0.020	−0.017
Autocorrelation 30: Inside	−0.008	0.006	−0.013	−0.012
Autocorrelation 300: Local	−0.008	0.000	−0.008	−0.007
Autocorrelation 300: Inside	−0.001	0.000	−0.001	−0.001
Variance Ratio 30/300: Local	−0.036	0.006	−0.043*	−0.032*
Variance Ratio 30/300: Inside	−0.037	0.006	−0.043**	−0.033**

Table 9. Impact on Cross Market Arbitrage Analysis and Trading Volume

Panel A summarizes the arbitrage opportunities arising on the Euronext order books during the pre-event period i.e., between 1 December 2008 and 13 January 2009. Section 5.3.4 describes how we identify each arbitrage opportunity. Frequency is the mean daily frequency of arbitrage opportunities between 08:01 and 16:29, Time Negative Spread is the total amount of time between 08:01 and 16:29 when the two order books are crossed, Arbitrage Activity is the average daily volume attributed to the resolution of the arbitrage opportunities, as a percentage of total trading volume on the two Euronext order books during the continuous trading sessions. Panel B reports the impact of introducing a single order book on the total Euronext trading volume. We estimate a difference-in-difference regression for the level and logarithm of the dependent variable, and report the coefficient of the interaction term which equals one for all treatment stocks and for all days on or after 19 January 2009, and zero otherwise. We employ stock and day fixed effects and double cluster standard errors by stock and day. *, **, *** denote significance at 10%, 5%, and 1%, respectively.

Panel A: Arbitrage Analysis

Stock	Arbitrage	Time	Arbitrage Activity	
	Frequency	Negative Spread	€ '000 Volume	Share
DEXI	155.8	10.3%	897	10.5%
FOR	219.3	7.5%	1,697	7.2%
ISPA	717.5	4.0%	11,086	6.1%
UNBP	173.5	2.9%	1,843	4.9%
GLPG	2.4	4.7%	8	4.3%
ONCOB	0.3	1.9%	0	0.3%
RCUS	1.3	7.7%	7	6.0%
VRKP	1.2	8.9%	6	14.2%
THEB	0.5	2.2%	2	15.3%
MEAN	141.3	5.6%	1,727,335	6.2%

Panel B: Trading Volume

	Treatment	Control	Effect Size	
	Post-Pre	Post-Pre	Levels	Logs
Total Volume	−137	−2,079	1,935	0.406

53

Table 10. Testing for Parallel Trends

This table describes the results from the difference-in-differences regression augmented by including interaction between the treatment variable and six weekly dummies on both sides of the event date. The weekly dummies W_k are defined based on calendar weeks with $k \in -6, -5, -4, \ldots, 4, 5, 6$. The week before the event date W_{-1} includes three additional days before the event date. We report the coefficients of these interaction terms. Additionally, we report the p-values from a Wald test for joint significance of the interaction between all weekly pre-event dummies and the treatment variable.

Variable	Effective Spread Local		Effective Spread Inside		Realized Spread 30 Local		Realized Spread 30 Inside		Realized Spread 60 Local		Realized Spread 60 Inside	
	Level	Log	Level	Log	Level	Log	Level	Log	Level	Log	Level	Log
W_{-6}	0.175	0.185	0.017	0.085	−0.427	0.417	−0.447	−0.083	−0.436	−0.091	−0.451	−0.153
W_{-5}	0.296	0.083	0.354	0.108	−0.666	−0.216	−0.389	0.112	−0.461	−0.106	−0.300	0.099
W_{-4}	0.269	0.088	−0.042	−0.005	−0.296	−0.045	−0.316	0.249	−0.270	−0.012	−0.267	0.226
W_{-3}	0.379	0.216	−0.320	0.035	−0.251	0.526***	−0.630	0.021	−0.487	0.210	−0.838	−0.389
W_{-2}	−0.233	0.235	−0.182	0.100	−0.666***	−0.125	−0.481	0.163	−0.427*	0.079	−0.235	0.343
W_{+1}	−0.717	−0.284	−0.706	−0.282	−1.119	0.151	−1.074	−0.097	−1.113	0.120	−1.068	−0.098
W_{+2}	−0.534	−0.302*	−0.541	−0.302*	−0.568	0.304	−0.479	−0.066	−0.574	0.315	−0.484	−0.053
W_{+3}	−0.745**	−0.325*	−0.734**	−0.324*	−0.481	0.132	−0.429	−0.143	−0.475	0.123	−0.423	−0.108
W_{+4}	0.044	0.052	0.058	0.054	−0.247	0.367	−0.180	−0.020	−0.240	0.384	−0.173	−0.011
W_{+5}	−0.182	−0.120	−0.161	−0.117	−0.469	−0.013	−0.472	0.079	−0.458	0.023	−0.461	0.106
$H0: W_k = 0 \,\forall\, k > 0$												
p-value	0.623	0.460	0.589	0.926	0.509	0.487	0.760	0.987	0.731	0.973	0.627	0.757

54

Recent Issues